DATA SLICER FOR COMBINED TRELLIS DECODING AND EQUALIZATION

[0001]   This application is a continuation-in-part of U.S. patent application number 09/099,730, filed June 19, 1998 and U.S. patent application number 09/100,705 filed June 19, 1998, the contents of which are incorporated herein by reference.

BACKGROUND OF THE INVENTION

5   [0002]   The present invention concerns data slicers or quantizers and in particular, a data slicer suitable for use in a baseband or passband decision-feedback equalizer.

[0003]   Equalizers are typically used in coded digital communications systems to compensate for multipath/linear filtering effects caused by the transmission channel. These effects are commonly referred to as channel impairments and include signal
10   distortion that may occur in the transmitter, in the receiver or in the channel through which the signal is transmitted. The equalizer is an adaptive filter, often implemented as a finite impulse response (FIR) filter, an infinite impulse response (IIR) filter or a combination of FIR and IIR filters. Each filter has a plurality of coefficients that are adapted to minimize an error criterion. This error criterion may be, for example, the
15   mean-square error between a transmitted training signal and the received training signal. A typical equalizer maintains a copy of the transmitted training signal to compare with the received training signal. It is generally believed that a decision-feedback equalizer (DFE) has better asymptotic performance than a linear equalizer as described in a text book by J. G. Proakis, entitled Digital Communications.

20   [0004]   A typical DFE is shown in Figure 1. The received signal is applied to an FIR filter and the output signal produced by the FIR filter is applied to an IIR filter. The IIR filter includes a subtracter 111, a slicer 112 and an IIR filter section 114. The subtracter 111 subtracts the filtered signal provided by the IIR filter section 114 from the output signal of the FIR filter 110. The slicer 112 quantizes the signal provided by the
25   subtracter 111 to produce an approximation of the signal that was transmitted. The IIR filter section, which may be, for example, an FIR filter in a feedback loop, processes the quantized signal to produce the signal that is subtracted by the subtracter 111. For an

uncoded modulation scheme, the DFE uses the slicer to get decisions for the feedback portion. The slicer is a nearest element decision device which returns the source constellation member closest to its input. In decision directed operation, the output signal of the slicer is compared to its input signal to determine in what way the

5   coefficients of the FIR and IIR filter sections should be updated to minimize any differences between the signal that was recovered and the known training signal.

[0005]   For a coded modulation scheme, it may be desirable to replace the slicer with a decoder, that may include, for example, a trellis decoder, a deinterleaver, and a Reed Solomon (RS) decoder. The use of such a decoder, however, delays the decision

10   on what symbol was transmitted by several symbol periods. These delays can be prohibitive for the DFE, since it relies on canceling the inter-symbol-interference of the previous symbols on the current symbol by using previously available decisions. Hence, the state of the art has typically not used a complete decoder, but a range of simplified decoders including the simple slicer 112, that does not perform any decoding. A typical

15   problem with using only a slicer in a DFE is a loss in performance due to incorrect decisions. Because an incorrect decision used in the DFE to remove inter-symbol interference (ISI) can cause further errors, this performance loss is known as "error propagation."

[0006]   More complex decoding techniques may also be used, for example,

20   Reduced-State Sequence Estimation (RSSE) and parallel decision feedback decoding (PDFD). These techniques are described in an article by V. Eyuboglu and S. Qureshi, entitled "Reduced-State Sequence Estimation for Coded Modulation on Intersymbol Interference Channels" IEEE Journal on Selected Areas of Communications, Aug. 1989. Furthermore, US Patent 5,056,117 entitled DECISION-FEEDBACK

25   EQUALIZATION WITH TRELLIS CODING to R. Gitlin, describes a method by which multiple possible decisions are fed back and the best among them is chosen using a given criteria. Other techniques are described in an article by A. Duel-Hallen and C. Heegard, entitled "Delayed Decision-Feedback Equalization", IEEE Transactions on Communications May 1989. All of the above cited references are incorporated herein

30   by reference for their teachings on combined equalization and decoding techniques.

[0007]   Generally the common idea among these decoders is to use multiple possible decisions or to use a more complicated trellis decoder that includes a channel state estimate. The implementation complexity of these approaches, however, is significant and may undesirably add to the cost of the decoder.

[0008]   U.S. patent no. 5,923,711 entitled SLICE PREDICTOR FOR A SIGNAL RECEIVER to Willming describes a slicer for a trellis coded vestigial sideband (VSB) signal that estimates a current symbol using a partial estimate of the current symbol which is derived from the previous symbol. The partial estimate is derived only from the previous symbol and is used to reduce the probability of error in the estimate of the current symbol. The system disclosed by Willming recovers carrier frequency from the pilot signal of the VSB signal. The pilot signal of a VSB signal, is not a good reference, however, as it may be corrupted both in frequency and phase by multipath distortion. In addition, the system disclosed by Willming may not operate properly when used with VSB signals that do not have pilot signals.

SUMMARY OF THE INVENTION

[0009]   The present invention is embodied in a quantizer that may be used to recover N-bit symbols from successive channel impaired input samples representing a trellis encoded signal. The quantizer includes a partial trellis decoder that traces back at least one symbol to generate an estimate of a subset of the N bits of each symbol and, based on the estimate, selects a decision device for a reduced constellation to generate an estimate of the N bits.

[0010]   According to another aspect of the invention, the trellis encoder that generates the input samples is based on a set-partitioned code with feedback convolutional encoding and the partial trellis decoder computes path metrics for the current symbol based on path metrics of the previous symbol. The best path metric of the previous symbol is used to select the reduced-constellation decision device that produces the current symbol.

[0011]   According to yet another aspect of the invention, the trellis encoder that generated the input samples is based on a set-partitioned code and the partial trellis

decoder computes path metrics for the current symbol based on path metrics of the previous symbol. The best path metric from among the path metrics of the current symbol, corresponding to a subset of states for the previous symbol, is used to select the reduced-constellation decision device that produces the current symbol.

5   [0012]   According to another aspect of the invention, the quantizer is used in an equalizer for a modulated digital signal. The equalizer includes a first adaptive filter that processes passband signals and a second adaptive filter that processes passband signals.

[0013]   According to yet another aspect of the invention, the quantizer is used in an equalizer for a modulated digital signal. The equalizer includes a first adaptive filter that processes passband signals, a demodulator and a second adaptive filter that processes baseband signals.

[0014]   According to another aspect of the invention, the quantizer is used in an equalizer for a modulated digital signal. The equalizer includes a first adaptive filter that processes baseband signals, a demodulator and a second adaptive filter that processes baseband signals.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015]   Figure 1 (prior art) is a block diagram of a decision feedback equalizer.

[0016]   Figure 2 is a block diagram of a digital signal receiver and decoder which includes an embodiment of the invention.

[0017]   Figure 3A is a block diagram of an exemplary equalizer suitable for use in the digital signal receiver shown in Figure 2.

[0018]   Figure 3B is a block diagram of an alternative equalizer suitable for use in the digital signal receiver shown in Figure 2.

[0019]   Figure 4 (prior art) is a block diagram of an encoder modulator for trellis coded modulation.

[0020]   Figure 5 (prior art) is a code constellation decomposition diagram which is useful for describing the operation of the encoded modulator shown in Figure 2.

5   **[0021]**   Figure 6 (prior art) is a block diagram of a generalized convolutional encoder structure with feedback.

[0022]   Figures 7A, 7B and 7C (all prior art) are block diagrams of exemplary specific convolutional encoders with feedback.

[0023]   Figure 8 (prior art) is a block diagram of an exemplary trellis encoder as
10   specified for digital television signals in the standard specified by the Advanced Television System Committee (ATSC).

[0024]   Figure 9A (prior art) is a constellation diagram for the trellis encoder shown in Figure 6.

[0025]   Figure 9B (prior art) is a trellis diagram for the trellis encoder shown in
15   Figure 6

[0026]   Figure 10 is a block diagram of a first exemplary symbol slicer according to the present invention.

[0027]   Figure 11 is a block diagram of a second exemplary symbol slicer according to the present invention.

20   **[0028]**   Figure 12 is a block diagram of a third exemplary symbol slicer according to the present invention.

DETAILED DESCRIPTION

[0029]   An exemplary receiver architecture is illustrated in Figure 2. A radio frequency (RF) demodulator 210 receives an RF signal that has been modulated using VSB and translates the frequency spectrum of the received signal to near baseband, so that its center frequency is approximately zero. The exemplary embodiment of the invention may also be used with QAM signals, as described in copending patent application no. 09/100,705 to Strolle et al. filed June 19, 1998. For the sake of simplicity, only the VSB demodulator is described below.

[0030]   The output signal of the demodulator 210 is not a demodulated baseband signal but a modulated passband signal. The exemplary RF demodulator is a synchronous demodulator that provides both in-phase (I) and quadrature-phase (Q) signal components. In the drawing Figures, the I and Q signal components are shown as a single line even though they may be conveyed by two conduction paths.

[0031]   The RF demodulator 210 also generates digital samples of the RF signal. These samples may be taken at a sample rate greater than the symbol rate (baud rate) of the modulated signal. In the exemplary embodiment of the invention, the samples provided by the RF demodulator 210 are at approximately four-times the symbol rate.

[0032]   The passband signal provided by the RF demodulator 210 is first interpolated and then filtered by a matched filter 214 before being sampled by the baud timing recovery circuit 212. The frequency spectrum of the matched filter 214 depends on how the signals were filtered at the transmitter. If, for example, the transmitter (not shown) applied a square-root raised cosine frequency response characteristic to the digital signal, the matched filter 214 also applies a square-root raised cosine characteristic. The cascade combination of these two filters provides a signal having minimal inter-symbol interference (ISI).

[0033]   Baud synchronization by the baud timing circuitry 212 is accomplished without knowing the exact carrier frequency or phase by band-edge, phase-lock techniques. These techniques are described in U.S. Patent numbers 5,872,815 entitled APPARATUS FOR GENERATING TIMING SIGNALS FOR A DIGITAL

TELEVISION RECEIVER and 5,799,037 entitled RECEIVER CAPABLE OF DEMODULATING MULTIPLE MODULATION DIGITAL FORMATS that are incorporated herein by reference for their teaching on baud synchronization. This method introduces minimal additional error under perfect signaling conditions.

5   **[0034]**   Next, the VSB signal is passed through the VSB pilot loop 216. This loop synchronizes a local oscillator signal to the pilot component of the VSB signal. This pilot signal is described in the above-referenced ATSC Digital Television Standard.

[0035]   The output signal of the VSB pilot loop 216 is applied to an equalizer/carrier loop recovery circuit 220, described below, that includes an
10   embodiment of the present invention. This circuitry compensates for multipath distortion and for the frequency response characteristic of the transmission channel. The loop recovery circuit also includes a phase locked loop (PLL) that determines the frequency and phase of the residual carrier signal to provide samples of the completely demodulated signal for application to the decoder 222. The decoder 222 may include a
15   VSB decoder.

[0036]   An exemplary adaptive equalizer structure is shown in Figure 3A. This equalizer includes a forward or finite impulse response (FIR) section 316 and a feedback or infinite impulse response (IIR) section including a summing circuit (e.g.. subtracter 318) and IIR filter element 330. Although the exemplary embodiment shows both an
20   FIR filter and an IIR filter, it is contemplated that an equalizer according to the subject invention may be made using only the FIR filter section 316 or only the IIR filter section. These alternate configurations are illustrated by the connections 317 and 319, shown in phantom, in Figures 3A and 3B. If the equalizer includes only the IIR section, then the FIR filter 316 is replaced by the connection 317. If only the FIR filter 316 is
25   used, then the subtracter 318 and filter element 330 are removed and the output of the equalizer is the signal y as indicated by the connector 319 or the output symbols provided by the slicer 324, as described below.

[0037]   In the exemplary embodiment of the invention shown in Figure 3A, demodulated QAM signals from the timing loop are applied to the FIR filter section
30   316. The FIR section receives data at a residual carrier frequency offset because, as

described above, the RF demodulator 210, shown in Figure 2, does not provide a baseband signal. Thus, the FIR filter 316 operates in the passband (away from DC). The FIR filter 316 is fully described at any given instant by its vector of tap weights (or coefficients) and written as a vector of impulse response coefficients $\mathbf{f} =$[f_0 f_1 ... f_{N-1}

5   $]^T$ where each of the values f_0 through f_{N-1} is a complex value. Similarly, the vector of tap weights for the feedback or IIR filter is fully described by a vector of impulse response coefficients $\mathbf{g} =$[g_0 g_1 ... g_{L-1} $]^T$, where each of the values g_0 through g_{L-1} is a complex value. The input signal to the FIR filter is the vector $\mathbf{r}$ and the output signal of the FIR filter is the scalar $y=\mathbf{r}^T\mathbf{f}$. The input signal to the IIR filter is the vector

10 $\mathbf{z}$ and the output signal of the IIR filter is the scalar $w=\mathbf{z}^T\mathbf{g}$. The output signal of the IIR filter is subtracted from the signal y in the subtracter 318 to produce the signal x. It is this signal x which is de-rotated and applied to the carrier loop circuitry 323.

[0038]   The carrier loop circuitry 323 generates a carrier signal corresponding to the carrier frequency of the passband signal provided by the RF demodulator 210

15 (shown in Figure 2). This circuitry does not rely on the pilot signal because, as described above, the frequency and phase of the pilot signal may, in some VSB signals, be severely attenuated or it may be corrupted due to multipath distortion. The carrier signal recovered by the carrier loop circuitry 323 is used to exactly demodulate the VSB signal to baseband in order to recover the data values. Exemplary carrier loop circuitry

20 is described in the above referenced patent application to Strolle et al. The equalizer and carrier loop circuitry shown in Figure 3A operate in two modes, acquisition and tracking. When the circuitry is operated in acquisition mode, the switch S4 connects the output port of the subtracter 318 directly to the input port of the IIR filter element 330. The equalizer is configured as a linear structure.

25 [0039]   When the circuitry is used for tracking, switch S4 is switched to add the slicer 324 into the feedback loop. The equalizer is configured as a nonlinear DFE structure. In the exemplary embodiment of the invention, the slicer 324 is part of carrier loop circuitry 323. In this configuration, the output signal of the IIR filter element 330 is subtracted from the output signal of the FIR filter 316 and then shifted to baseband by

30 multiplying it, in the mixer 322, by $e^{j\theta_k}$ the current estimate of the complex conjugate of the residual carrier offset. The baseband signal is then quantized by the slicer 324 in order to form the error calculations used to update the equalizer coefficients, as

described in the above-referenced patent application to Strolle et al. The baseband

signal is shifted back into the passband by multiplication, in the mixer 326 by $e^{-j\theta_k}$,

the current estimate of the residual carrier offset. The output signal of the mixer 326 is then fed back into the input port of the IIR filter element 330 via switch S4. The input signals to the carrier loop circuitry 323 are the input signal to the slicer 324 and the output signal of the slicer 324.

[0040]   The embodiments of the invention described below use a VSB slicer that includes a partial trellis decoder.

[0041]   Figure 3B is an alternative embodiment of the adaptive equalizer structure. In the embodiment shown in Figure 3B, the FIR filter section 316 processes passband signals while the IIR filter section, including the subtracter 318 and the IIR filter element 330, processes baseband signals. In Figure 3B, elements 310 and 316 operate in the same way as shown in Figure 3A. In the equalizer shown in Figure 3B, however, the passband signal provided by the FIR filter 316 is converted to a baseband signal by

mixer 320, that multiplies the passband signal by $e^{j\theta_k}$, the current estimate of the complex conjugate of the residual carrier signal. The output signal of mixer 320 is then applied to the subtracter 318, which subtracts the output signal of the IIR element 330 from the output signal of the mixer 320. The output signal of the subtracter 318 is then applied directly to one pole of switch S4 and through slicer 324 to the other pole of switch S4. When operating in acquisition mode, switch S4 applies the output signal of the subtracter 318 to the input port of the IIR filter element 330. When operating in tracking mode, switch S4 applies the output signal of the slicer 324 to the input port of the IIR filter element 330. As with the embodiment shown in Figure 3A, the input signals to the carrier loop circuitry 323 are the input signal to the slicer 324 and the output signal of the slicer 324.

[0042]   The symbol values returned by the slicers 324 are a key feature of the equalizers shown in Figures 3A and 3B. In these exemplary embodiments, the slicers 324 are being used for equalizer filtering and adaptation as well as for carrier phase tracking.

[0043]   To determine how a code may be effectively decoded, it is desirable to study the codes currently used in the state of the art. Generally the state of the art uses "Set -Partition Codes" or "Coset Codes", and will be referred to as set-partition codes in the following. These codes use a feedback convolution code and a subset mapping

5   technique, as described in an article by G. Ungerboeck, entitled, "Trellis-Coded Modulation with Redundant Signal Sets Part I and II", IEEE Transactions on Communications, Feb. 1987, which is incorporated herein by reference for its teaching on set partition codes.

[0044]   An example of a set partition encoder is shown in Figure 4, which is

10   copied from the above-referenced article by Ungerboeck. The coded bits select subsets of a specific constellation as shown, for example, in Figure 5. The uncoded bits then select individual points within the constellation. Note that the distance between signal points increases by a factor of the square root of 2 as shown in Figure 4. It is noted that the description of distances in this article are erroneous. Nonetheless, if the coded bits

15   can somehow be 'predicted' then an appropriate slicer can be selected, and because of the greater inter-symbol distance in the constellations, this slicer will provide better decisions than a simple slicer.

[0045]   The method of predicting the coded bits relies on a second property of the set-partition codes: the use of a feedback convolutional encoders with a certain

20   characteristic to generate the codes. The class of these feedback encoders includes those that are used to maximize the minimum Euclidean distance as described in an article by G. Ungerboeck, entitled, "Channel Coding With Multilevel/Phase Signals", IEEE Transactions on Information Theory, Jan. 1982, which is incorporated herein by reference for its teaching on set partition codes. This class of feedback encoders has the

25   property that some coded bits (at least the last coded bit shown as $Z0(n)$) are determined uniquely one interval before the bits are used. This is illustrated by Figures 6 and 7A through 7C. Note that in Figures 7A through 7C, the 'coded bit' $y0(n)$ is determined one symbol interval before the current symbol. Hence, the previous state estimate can be used to determine $y0(n+1)$. Clearly, this observation can be generalized to the case

30   of many coded bits being estimated.

[0046]   Many methods can now be used to generate an estimate of the coded bits. Generally, a trellis decoder uses a Viterbi Algorithm (VA) with path length >> 2, as described in a text book by A. Viterbi et al. entitled Principles of Digital Communications and Coding which is incorporated herein by reference for its teaching
5   on trellis encoding and decoding.

[0047]   The present invention is embodied in a slicer that provides a better estimate for the feedback portion of the DFE than a slicer which uses no coding information. In addition, the complexity of the trellis decoder and the decision-feedback portion is not increased. Furthermore, no delay is incurred by use of the slicer. Hence,
10   this approach provides a better estimate of the symbol than the nearest element decision device with no delay compared to a Viterbi decoder.

[0048]   The slicer may be used with a class of codes generally referred to as " Set Partition Codes" or " Coset Codes." The slicer relies on two properties of these codes: the use of a set-partitioned signal mapper and the use of a feedback convolutional code.
15   The slicer generates an estimate of some of the coded bits and then uses this estimate to select a slicer from among a set of slicers. Several exemplary methods are described to generate an estimate of the coded bits.

[0049]   The most complex of these methods is to use a complete trellis decoder with path length = 2. In this case, the trellis decoder computations are desirably finished
20   to generate the estimate of the coded bits, then the coded bits are used to generate the output signal of the slicer. This output signal is then used in the feedback filter of the DFE for filtering and adaptation. These steps imply, however, that both the output signal of the feedback filter and the trellis decoder computation are completed within one symbol interval.

25   [0050]   The simplest of the described methods for generating an estimate of the coded bits is to use a complete trellis decoder with path length = 1. In this case, the computations performed by the feedback filter and the trellis decoder are done in parallel, i.e., the estimate of the coded bits is done before the next signal sample is available. Hence the slicer selection for the current symbol is implemented before the
30   current symbol is available. This implementation is very simple.

[0051]   Another method of generating an estimate of the coded bits is to use only a subset of states for the selection of the coded bits, but to keep track of the full trellis decoder after the estimation of the coded bits is complete. Using this method, the computations performed during a symbol interval include first computing an estimate of coded bits, then in parallel, using the output signal of the slicer to compute the feedback portion of the DFE, thus completing the trellis decoder computation. This method can use ideas from the M-algorithms, described in the above-referenced book by Viterbi et al., and selects the best M states to make the estimation of the coded bits.

[0052]   Finally, a simplified mechanism of generating an estimate of the coded bit for the VSB system is described, which method is specified relative to the ATSC trellis encoder.

[0053]   For the purpose of describing the invention, the VSB trellis code as used in the ATSC system is used as an example. It is contemplated, however, that the subject invention has general application to other systems which use trellis codes or other Viterbi-type encoding techniques.

[0054]   The ATSC trellis encoder shown in Fig. 8 defines states $S_1(n)$ and $S_0(n)$ as the current input samples to the delay elements. Note that, by operation of the delay elements 810 and 812, the trellis encoder shown in Fig. 8 is constrained by equations (1) and (2).

$$Z_0(n) = S_0(n-1) \tag{1}$$

$$Z_0(n) = S_1(n) \tag{2}$$

[0055]   The trellis code output $(R(n))$ is completely specified by the previous state $(S_1(n-1), S_0(n-1))$ and the current inputs $Z_1(n)$ and $Z_2(n)$. Specifically, the bit $Z_0(n)$ can be specified by the previous states $(S_1(n-1), S_0(n-1))$ and the current input $Z_1(n)$.

[0056]   Such a relationship may be described using the constellation 900 shown in Figure 9A and the corresponding trellis diagram shown in Fig. 9B. In Figure 9B, The branches specify a transition from a specific state $(S_1(n-1), S_0(n-1))$ of the trellis

encoder to another state $(S_1(n), S_0(n))$ given the input bit $Z_1(n)$. A broken line branch represents a transition when $Z_1(n)=0$, and a solid line represents a transition when $Z_1(n)=1$. Also listed on top of the transitions are one of the four subsets 'a', 'b', 'c'. or 'd' that are defined as follows.

5

'a':$\{-7,1\}$ (corresponds to $Z_1(n)=0$ and $Z_0(n)=0$)

'b':$\{-3,5\}$ (corresponds to $Z_1(n)=1$ and $Z_0(n)=0$)

'c':$\{-5,3\}$ (corresponds to $Z_1(n)=0$ and $Z_0(n)=1$)

'd':$\{-1,7\}$ (corresponds to $Z_1(n)=1$ and $Z_0(n)=1$)

[0057]   Thus the trellis diagram of Fig. 9B completely specifies the next state as

10   well as the next output given the new input to the state diagram. Note that the bit $Z_2(n)$ selects one of two signal points from subsets 'a', 'b', 'c', or 'd'.

[0058]   The receiver shown in Figures 2 and 3A or 3B uses a slicer according to the present invention. This slicer provides an estimate $Z_0(n)$. Exemplary slicers for VSB are shown in Figures 10 and 11. A first exemplary implementation estimates a

15   single bit, Z_0, of the three bit symbol. Figure 10 shows an exemplary implementation of a slicer 324 (shown in Figures 3A and 3B) when bit Z_0 is known. As shown in this implementation the 8 VSB constellation may be decomposed into the disjoint union of two 4 VSB constellations. In this exemplary embodiment, a trellis decoder 1010 may be used to estimate the one bit Z_0, and this value may then be used to control the

20   multiplexer 1012 to select the symbols from one of two disjoint 4 VSB quantizers 1014 and 1016. For example, the 8 VSB constellation can be decomposed as set forth in equation (3).

$$8 \text{ VSB} = \{-7,-3,1,5\} \cup \{-5,-1,3,7\} \tag{3}$$

[0059]   where the first 4 VSB subset, (quantizer 1014) corresponds to $Z_0=0$ and

25   the second VSB subset (quantizer 1016) corresponds to $Z_0=1$ respectively.

[0060] While a general trellis decoder 1010 can be used to estimate Z_0, it is desirable that the estimate for Z_0 be available within one symbol interval. The decoder becomes more complicated if the estimate is not available within one symbol interval because inter-symbol interference (ISI) cancellation is less reliable. First, the decision feedback cannot be done within one symbol interval. Assuming that the Z_0 estimate is available N trellis encoded symbols later (e.g. because the ATSC standard specifies 12 parallel trellis encoders, the decision feedback cannot be done for 12*N symbols periods). Furthermore, the computation of the error term is also delayed. Consequently, the data values are stored to implement an LMS-like adaptive algorithm. The longer the delay, the larger the storage area to hold the data values. Hence, it is desirable to get an estimate of Z_0 in the same symbol interval in which the value including Z_0 is decoded.

[0061] In general, trellis decoders are well known in the art and use the Viterbi Algorithm as described in the above-referenced text by Viterbi et al. During every interval, path metrics corresponding to every state are computed in accordance with a trellis diagram such as that shown in Figure 9. The path metrics are defined as $p_0(n)$, $p_1(n)$, $p_2(n)$, $p_3(n)$ for the four-state trellis corresponding to states $S_1(n)$, $S_0(n)=(0,0)$, $(0,1)$, $(1,0)$, $(1,1)$ respectively. The branch metrics, bm, for the branches in the trellis are denoted as shown in equations (4) - (7).

$$bm(\text{'a'}) = \text{Euclidean distance (current sample, subset ('a'))} \quad (4)$$

$$bm(\text{'b'}) = \text{Euclidean distance (current sample, subset ('b'))} \quad (5)$$

$$bm(\text{'c'}) = \text{Euclidean distance (current sample, subset ('c'))} \quad (6)$$

$$bm(\text{'d'}) = \text{Euclidean distance (current sample, subset ('d'))} \quad (7)$$

[0062] Then the path metrics are computed every symbol interval according to the Viterbi Algorithm as shown in equations (8) - (11).

$$p_0(n) = \text{minimum } \{p_0(n-1)+bm(\text{'a'}), p_2(n-1)+bm(\text{'b'})\} \quad (8)$$

$$p_1(n) = \text{minimum } \{p_0(n-1)+bm(\text{'b'}), p_2(n-1)+bm(\text{'a'})\} \quad (9)$$

$$p_2(n) = \text{minimum } \{p_1(n-1)+bm(\text{`c'}), p_3(n-1)+bm(\text{`d'})\} \qquad (10)$$

$$p_3(n) = \text{minimum } \{p_1(n-1)+bm(\text{`d'}), p_3(n-1)+bm(\text{`c'})\} \qquad (11)$$

[0063]   This notation is used below to describe the different methods for estimating $Z_0(n)$.

Case 1: Path Memory = 2

[0064]   In this case, the $Z_0(n)$ estimate is based on path metrics $p_0(n)$, $p_1(n)$, $p_2(n)$, $p_3(n)$. If the minimum path metric, p, among $\{ p_0(n), p_1(n), p_2(n), p_3(n) \}$ corresponds to either states 0 or 1 (i.e. $p_0(n)$ or $p_1(n)$) then the estimate for $Z_0(n)$ equals 0, otherwise the estimate equals 1. The rationale for this selection is based on equation (2) above. This case may be described by equation (12).

$$Z_0(n)=0 \qquad \text{if minimum } \{p_0(n), p_1(n)\} < \text{minimum } \{p_2(n), p_3(n)\} \quad (12)$$

$$=1 \qquad \text{otherwise}$$

Case 2: Path Memory = 1

[0065]   In this case, the $Z_0(n)$ estimate is based on path metrics $p_0(n-1)$, $p_1(n-1)$, $p_2(n-1)$, $p_3(n-1)$. If the minimum path metric among $\{p_0(n-1), p_1(n-1), p_2(n-1), p_3(n-1)\}$ corresponds to either states 0 or 2, then the estimate for $Z_0(n)$ equals 0, otherwise the estimate equals 1. The rationale for this selection is based on equation (1) above. This case may be described by equation (13)

$$Z_0(n)=0 \quad \text{if minimum } \{p_0(n-1), p_2(n-1)\}<\text{minimum } \{p_1(n-1), p_3(n-1)\} \quad (13)$$

$$=1 \qquad \text{otherwise.}$$

[0066]   Clearly the advantage of this case over path memory = 2 case is that the computation of the estimate for $Z_0(n)$ can be made before the signal sample is available.

Case 3: Path Memory = 2, (low complexity)

[0067] In this case, the estimate of $Z_0(n)$ is based on path metrics $\{p_0(n-1), p_1(n-1), p_2(n-1), p_3(n-1)\}$, and on 'extending' the two best states corresponding to the minimum two path metrics from among $\{p_0(n-1), p_1(n-1), p_2(n-1), p_3(n-1)\}$. In this case, the possible results are given by the conditional statements shown in Tables 1 and 2.

Table 1

If $p_0(n-1)$ and $p_2(n-1)$ are the best states then $Z_0(n)=0$
If $p_1(n-1)$ and $p_3(n-1)$ are the best states then $Z_0(n)=1$

Table 2

If $p_0(n-1)$ and $p_1(n-1)$ are the best states then
  if minimum $\{p_0(n-1) + bm('a'), p_0(n-1) + bm('b')\}$
    $< $ minimum $\{p_1(n-1) + bm('c'), p_1(n-1) + bm('d')\}$ then $Z_0(n) = 0$
  else $Z_0(n) = 1$

[0068] It is possible to simplify the final computation in Table 2 to:

   if $p_0(n-1) + bm('ab') < p_1(n-1) + bm('cd')$ then $Z_0(n) = 0$

   else $Z_0(n) = 1$

[0069] Where $bm('ab')$ is given by equation (14) and $bm('cd')$ is given by equation (15).

$$bm('ab') = \text{Euclidean distance (current sample, } Z_0=0 \text{ slicer)}$$
$$= \text{Euclidean distance(current sample, } \{-7,-3,1,5\}) \qquad (14)$$

$$bm('cd') = \text{Euclidean distance (current sample, } Z_0=1 \text{ slicer)}$$
$$= \text{Euclidean distance (current sample, } \{-5,-1,3,7\}) \qquad (15)$$

[0070]   The expression in Table 2 may be used for the rest of the pairs of path metrics $\{p_0(n-1),p_3(n-1)\}$, $\{p_2(n-1),p_1(n-1)\}$, and $\{p_2(n-1), p_3(n-1)\}$ by substituting $p_2(n-1)$ for $p_0(n-1)$ and by substituting $p_3(n-1)$ for $p_1(n-1)$ wherever appropriate.

[0071]   Clearly in this case, only one comparison is implemented to determine the $Z_0(n)$ estimate and thus this case is of lower complexity than case 1.

[0072]   Given these methods of estimating $Z_0(n)$, the input sample is quantified from one of two 4 VSB subsets. The symbol error for 4 VSB at the SNR corresponding to the ATSC threshold of visibility (TOV) is well below the rule of thumb of one in ten. Hence, the smart-slicer can provide reliable symbol quantization for decision-directed operation even at the ATSC specified TOV.

[0073]   So far the exemplary methods have calculated only the estimate of the bit $Z_0(n)$. It is possible, however, to estimate the bit $Z_1(n)$ as well. In this option, the bits $Z_0(n)$ and $Z_1(n)$ are estimated and the 8 VSB constellation may be decomposed into a disjoint union of four 2 VSB subsets as set forth in equation (16)

$$8 \text{ VSB} = \{-7,1\} \cup \{-5,3\} \cup \{-3,5\} \cup \{-1,7\} = \text{`a'} \cup \text{`c'} \cup \text{`b'} \cup \text{`d'} \quad (16)$$

where the four 2 VSB subsets correspond to $(Z_1,Z_0) = (0,0), (0,1), (1,0), (1,1)$ respectively. An implementation of this slicer 112''is shown in Figure 11, where a path length = 2 trellis decoder 1110 generates estimates of Z_0 and Z_1 and applies these estimates as control signals to a multiplexer 1112. The multiplexer 1112 divides the symbols among quantizers 1114, 1116, 1118, and 1120 according to the values of Z_0 and Z_1.

[0074]   The bits Z_0 and Z_1 may be estimated using the conditional statements of Table 1 shown above for Path Memory = 2. Given $p_0(n),p_1(n)$, $p_2(n)$, $p_3(n)$, the process calculates the minimum path metric among $\{p_0(n), p_1(n), p_2(n), p_3(n)\}$, then determines whether the path that was selected for the minimum path metric corresponds to either $Z_1(n)=0$ or $Z_1(n)=1$ (i.e., either the dotted or the solid line in Figure 7B). The $Z_0(n)$ estimate may be calculated as described above. Hence, using the method outlined above, both the $Z_0(n)$ and $Z_1(n)$ estimates are available.

[0075]   For the reduced complexity Case 2 described above, a further determination may be made of whether the path selected for the minimum path metric corresponding to either $Z_1(n)=0$ or 1 (as described above) may be used as an estimate of $Z_1(n)$.

[0076]   While this technique has been described only for the ATSC trellis code, it is clear from the above discussion that it can be applied for any feedback trellis encoder that may be partitioned into distinct sets. In particular, although the discussion set forth above describes a three-bit trellis code, it is contemplated that the same method may be applied to trellis codes that represent digital values having four or more bits. In addition, although the exemplary embodiments of the invention generate estimates for one or two bits to decode the remaining bits, it is contemplated that more bits may be estimated using extensions of the apparatus and method described above.

[0077]   In general, the techniques described above may be used to implement a decoder of the type shown in Figure 12. In this Figure, the code is divided into N subsets, each corresponding to one of the trellis decoders 1210a through 1210n. input samples are applied to the decoder so that each of the trellis decoders 1210a through 1210n acts on the input samples in parallel. At the same time, the samples are processed by a bit prediction/multiplexer control processor 1212, which, using the techniques described above, generates an estimate for a subset of the bits of the digital values that defines which one of the trellis decoders 1210a through 1210n is most likely to produce the correct result. The output signal of the processor 1212 is applied to a multiplexer 1214 that selects the output signal of the appropriate decoder 1210a through 1210n as the symbol estimate.

[0078]   In the embodiment shown in Figure 12, the bit prediction and multiplexer control processor 1214 may implement any of the bit estimation techniques described above.

[0079]   While the invention has been described in terms of exemplary embodiments, it may be practiced as outlined above within the scope of the appended claims.